arXiv:1607.03954v1 [stat.ME] 13 Jul 2016

Ensemble preconditioning for Markov chain Monte Carlo simulation

Charles Matthews*
Department of Statistics
University of Chicago
Chicago IL, 60615
c.matthews@uchicago.edu

Jonathan Weare
Department of Statistics and
James Frank Institute
University of Chicago
Chicago IL, 60615
weare@uchicago.edu

Benedict Leimkuhler
School of Mathematics
University of Edinburgh
Edinburgh, EH93FD, UK
b.leimkuhler@ed.ac.uk

July 15, 2016

Abstract

We describe parallel Markov chain Monte Carlo methods that propagate a collective ensemble of paths, with local covariance information calculated from neighboring replicas. The use of collective dynamics eliminates multiplicative noise and stabilizes the dynamics thus providing a practical approach to difficult anisotropic sampling problems in high dimensions. Numerical experiments with model problems demonstrate that dramatic potential speedups, compared to various alternative schemes, are attainable.

1 Introduction

A popular family of methods for Bayesian parameterization in data analytics are derived as Markov chain Monte Carlo (MCMC) methods, including Hamiltonian (or hybrid) Monte Carlo (HMC)[7], or the Metropolis adjusted Langevin algorithm (MALA)[27]. These methods involve proposals that are based on approximating a continuous time (stochastic) dynamics that exactly preserves the target (posterior) density, followed by an accept/reject step to correct for

*corresponding author

approximation errors. In some cases the Metropolis–Hastings correction may be omitted, provided the consequent approximation bias is well controlled. Efficient parameterization of the stochastic differential equations used in these procedures has the potential to greatly accelerate their convergence, particularly when the target density is poorly scaled, i.e., when the Hessian matrix of the logarithm of the density has a large condition number. In precise analogy with well established strategies in optimization, the solution to conditioning problems in the sampling context is to find a well chosen change of variables (preconditioning) of the system, such that the natural scales of the transformed system are roughly commensurate.

In this article we discuss an approach to dynamic preconditioning based on simultaneously evolving an ensemble of parallel MCMC simulations, each of which is referred to as a "walker." As we will show, the walkers provide information that can greatly improve the efficiency of MCMC methods. Our work is motivated by the method in [15] and earlier similar schemes that use multiple parallel simulations to improve proposal densities in MCMC simulations (see e.g. [10, 3, 5]). The authors of [15] highlight the notion of affine invariance which means, loosely, that the performance of the method when applied to sample the probability distribution with density π is identical to its performance when applied to sample $\pi_{A,v}(x) \propto \pi(Ax+v)$ regardless of the choice of invertible matrix A and vector v. Affine invariant methods are advantageous when sampling a severely anisotropic density, or when a suitable parameterization length scale for the problem is unknown. In this article, we however intentionally deviate from affine invariance to overcome practical limitations of earlier schemes. The methods in this paper also differ from those in [15] in that, like the methods listed in the opening paragraph, they require the first derivative of π, giving high accuracy even in the absence of a Metropolis–Hastings step, and making the methods scalable to high dimension.

Our starting point is the discrete approximation of a system of stochastic differential equations (SDEs),

$$\dot{x} = [J(x) + S(x)]\nabla \log \pi(x) + \operatorname{div}(J(x) + S(x)) + \sqrt{2S(x)}\,\eta(t) \qquad (1)$$

where $J(x)$ and $S(x)$ are skew-symmetric and symmetric positive semi-definite $N \times N$ matrices, respectively, with $\eta(t)$ representing a vector of independent Gaussian white noise components. In our sampling schemes, each walker generates a discrete time approximation of (1) with its own particular choice of J which corresponds to a notion of the localized and regularized sample covariance matrix across the ensemble of walkers and incorporates information about the target density π into the evolution of each walker.

Many optimization methods can be characterized as time discretizations of (1), with $J = 0$ and choices of S that encapsulate scaling information based upon higher derivatives of the model or a history of previously computed points. Modification of S to improve convergence also has been considered in the Monte Carlo literature, dating at least to [1]. This idea has been the focus of renewed attention in statistics and several new schemes of this or closely related types

have been proposed [22, 12]. From a more theoretical point of view, several authors have recently considered the effects of the choice of J and S on the ergodic properties of the solution (1) (see for example [25, 8, 19, 18]). In this paper we are concerned with motivating and presenting a particular choice of S and J that yields practical and efficient sampling schemes and we will leave a more detailed analysis to later work. Nonetheless, our development reveals the central role of discretization stability properties in determining the efficacy of a particular choice of J or S.

2 Preconditioning strategies for sampling

As in any MCMC scheme, the goal is to estimate the average $\mathrm{E}[f] = \int f(x)\pi(x)\mathrm{d}x$ by a trajectory average of the form

$$\overline{f}_N = \frac{1}{N}\sum_{n=0}^{n-1} f(x^{(n)}),$$

for large N. In many cases we can expect the error in an MCMC scheme to satisfy a central limit theorem: $\sqrt{N}\left(\overline{f}_N - \mathrm{E}[f]\right) \rightarrow_{\mathrm{dist}} N(0, \tau\sigma^2)$, where σ^2 is the variance of f under π (and is independent of the particular MCMC scheme), the quantity τ being the integrated auto-correlation time (IAT) which is often used to quantify the efficiency of an MCMC approach (see Appendix A).

To emphasize an analogy with optimization, for the moment assume that $J = 0$. The steepest descent algorithm of optimization corresponds to the so-called overdamped Langevin (or Brownian) dynamics,

$$x^{(n+1)} = x^{(n)} + \delta t\, \nabla \log(\pi(x^{(n)})) + \sqrt{2\delta t}\, \mathrm{R}^{(n)} \tag{2}$$

where here we have used an Euler-Maruyama discretization [24] and $\mathrm{R} \sim N(0, I)$. Discretization introduces an $O(\delta t)$ error in the sampled invariant distribution so a Metropolis–Hastings accept/reject test may be incorporated in order to recover the correct statistics (see the MALA algorithm [27]) when time discretization error dominates sampling error. Reducing δt gives a more accurate approximation of the evolution of the dynamics, and boosts the acceptance rate.

When π is Gaussian with covariance Σ, one can easily show that the cost to achieve a fixed accuracy depends on the condition number $\kappa = \lambda_{max}/\lambda_{min}$ where λ_{max} and λ_{min} are the largest and smallest eigenvalues of Σ. Indeed, one finds that the worst case IAT for the scheme in (2) over observables of the form $v^{\mathrm{T}}x$ is $\tau = \kappa - 1$ (see Appendix A). In this formula, the eigenvalue λ_{min} arises due to the discretization stability constraint on the step-size parameter δt and λ_{max} appears because the direction of the corresponding eigenvector is slowest to relax for the continuous time process. The presence of λ_{min} in this formula indicates that analysis of the continuous time scheme (1) (i.e. neglect of the discretization stability constraint) can be misleading when considering the effects of poor conditioning on sampling efficiency. Since the central limit theorem

suggests that the error after N steps of the scheme is roughly proportional to $\sqrt{\tau/N}$, the cost to achieve a fixed accuracy is again roughly proportional to κ.

Continuing to use $J = 0$ and taking, for example, $S(x) = -(\nabla^2 \log(\pi(x)))^{-1}$ in (1) and discretizing, we obtain a stochastic analogue of Newton's method,

$$x^{(n+1)} = x^{(n)} + \delta t\, S(x^{(n)}) \nabla \log(\pi(x^{(n)})) + \delta t\, \mathrm{div} S(x^{(n)}) + S^{1/2}(x^{(n)}) \sqrt{2\delta t}\, \mathrm{R}^{(n)}. \tag{3}$$

Schemes of similar form though neglecting the div$S$ term (and therefore requiring Metropolization) have been explored recently in [22]. Metropolization may also be used to correct the $\mathcal{O}(\delta t)$ sampling bias introduced by the discretization. It can be shown that the scheme is affine invariant in the sense that when it is applied to sampling $\pi_{A,v}$ it generates a sequence of samples $y^{(n)}$ so that $x^{(n)} = Ay^{(n)} + v$ has exactly the same distribution as the sequence of samples generated by the method when applied to π. We therefore expect that when this method can be applied (e.g. when the Hessian is positive definite), it should be effective on poorly scaled problems. This affine invariance property is shared by the deterministic Newton's method (obtained from (3) by dropping the noise term) and is responsible for its good performance when applied to optimize poorly scaled functions (e.g. when the condition number of the Hessian is large). We stress that the key to the usefulness of either the deterministic or stochastic Newton's method is that one does not need to make an explicit choice of the matrix A or the vector v. As the performance is independent of the choice of A and v, we can assume that A or v are chosen to remove poor scaling within the problem.

Due to the presence of the divergence term in the continuous dynamics, discretization will require evaluation of first, second and third order derivatives of $\log(\pi(x))$, making it prohibitively expensive for many models. To avoid this difficulty, one can estimate the divergence term using an extra evaluation of the Hessian (see Appendix F), or omit the divergence term and rely on a Metropolization step to ensure correct sampling. Regardless of how this term is handled, the system (3), unlike (2), is based on multiplicative noise which is known to introduce complexity (and reduce accuracy) in numerical discretization [24].

More fundamentally, complex sampling problems will exhibit regions of substantial probability where the Hessian fails to be positive definite. A choice that is both less costly and more robust is $S = \Sigma$, where we assume that Σ is the covariance matrix of π (even when π is not Gaussian) and is positive definite. It can be shown that the iteration in (3) is again affine invariant. The resulting scheme, which can be regarded as a simple quasi-Newton type approach, is closely related to adaptive MCMC approaches [26, 17]. On the other hand, because this choice of S does not depend on position, the scheme can be expected to perform poorly on problems for which the scaling behavior is dramatically different in different regions of space (e.g. the Hessian has high condition number and its eigenvectors are strongly position dependent), see Figure 1. This observation along with our comments on the method corresponding to $S(x) = -(\nabla^2 \log(\pi(x)))^{-1}$ suggest a choice of S corresponding to a notion of

(a) (b) (c)

Figure 1: We plot three examples of posterior distribution functions that might be described as poorly scaled. The distribution (a) has a scaling that can be removed through a linear change of variables, whereas useful scaling information in distributions (b) and (c) depends on the location in space. Proposals can benefit from taking into account local scaling behavior over the global covariance information (the condition number of their covariance matrices in both (b) and (c) are unity).

local covariance.

While a notion of local covariance will be central to the schemes we eventually introduce, we choose to incorporate that information not through S in (1), but through the skew-symmetric matrix J in that equation. In the remainder of this section we discuss how the choices of S described so far, and the corresponding properties of (3), have analogues in choices of J and a family of so-called underdamped Langevin schemes that we next introduce.

A popular way to obtain an MCMC scheme with decreased IAT relative to the overdamped scheme in (2) is to introduce "inertia." We extend the space by writing our state $x = (q, p)^{\mathrm{T}} \in \mathcal{D} \times \mathbb{R}^N \subset \mathbb{R}^{2N}$, with the target distribution $\hat{\pi}$ written as the product

$$\hat{\pi}(x) = \hat{\pi}(q, p) = \pi(q)\varphi(p), \qquad \int \hat{\pi}(q, p)\mathrm{d}p = \pi(q), \tag{4}$$

with the distribution of interest $\pi(q)$ recovered as the marginal distribution of the position vector q. For the distribution $\varphi(p)$ we will follow common practice and use $\varphi(p) \propto \exp(-\|p\|^2/2)$. With this extension of the space, we can obtain the standard underdamped form of Langevin dynamics for the choice

$$J = \begin{bmatrix} 0 & -I_N \\ I_N & 0 \end{bmatrix}, \qquad S = \begin{bmatrix} 0 & 0 \\ 0 & \gamma I_N \end{bmatrix} \tag{5}$$

in equation (1), where I_N is the $N \times N$ identity matrix and γ is a positive constant [24]. Previous work, especially in connection with molecular dynamics [21], has examined efficient ways to discretize the Langevin dynamics system while minimizing the error in sampling $\pi(q)$.

To incorporate information such as the Hessian matrix or the covariance matrix (or local covariance matrices) in the underdamped Langevin scheme, we focus on choices of J and S as follows:

$$J(x) = \begin{bmatrix} 0 & -B(q) \\ B(q)^{\mathrm{T}} & 0 \end{bmatrix}, \qquad S = \begin{bmatrix} 0 & 0 \\ 0 & \gamma I_N \end{bmatrix},$$

resulting in the system

$$\dot{q} = B(q)p, \qquad \dot{p} = B(q)^{\mathrm{T}}\nabla \log(\pi(q)) + \mathrm{div}(B(q)^{\mathrm{T}}) - \gamma p + \sqrt{2\gamma}\eta(t). \tag{6}$$

Discretization of the stochastic system may be derived by mimicking the BAOAB scheme [21]. Given a stepsize $\delta t > 0$, define $\alpha = \exp(-\gamma\delta t)$ and approximate the step from t_n to $t_{n+1} = t_n + \delta t$ by the formulas

$$p^{(n+1/2)} = p^{(n)} + \frac{\delta t}{2}F(q^{(n)}), \qquad q^{(n+1/2)} = q^{(n)} + \frac{\delta t}{2}B(q^{(n+1/2)})p^{(n+1/2)} \tag{7a}$$

$$\hat{p}^{(n+1/2)} = \alpha p^{(n+1/2)} + \sqrt{1-\alpha^2}\mathrm{R}^{(n)} + \frac{(\alpha+1)\delta t}{2}\mathrm{div}(B(q^{(n+1/2)})^{\mathrm{T}}) \tag{7b}$$

$$p^{(n+1)} = \hat{p}^{(n+1/2)} + \frac{\delta t}{2}F(q^{(n+1)}), \qquad q^{(n+1)} = q^{(n+1/2)} + \frac{\delta t}{2}B(q^{(n+1/2)})\hat{p}^{(n+1/2)} \tag{7c}$$

where $\mathrm{R} \sim N(0, I_N)$ and $F(q) = B(q)^{\mathrm{T}}\nabla \log \pi(q)$.The choice of matrix BB^{T} introduced in the next section will be a sum of the identity and a small (relative to the dimension N) number of rank 1 matrices, alleviating storage demands and reducing the cost of all calculations involving B to linear in N. As described in Appendix B, schemes of the form in (7) can also be used to generate proposals in a Metropolis–Hastings framework to strictly enforce a condition that, like detailed balance, guarantees that π is exactly preserved.

Suppose that, when applied to sampling the density $\pi_{A,v}$, an underdamped Langevin scheme of the form in (7) generates a sequence $(q^{(n)}, p^{(n)})$. The scheme will be referred to as affine invariant if the transformed sequence $(Aq^{(n)} + v, p^{(n)})$ has the same distribution as the sequence generated by the method when applied to sample π. As for (3) one can demonstrate that the choices $B(q)B^{\mathrm{T}}(q) = -(\nabla^2 \log(\pi(q)))^{-1}$ and $B(q)B^{\mathrm{T}}(q) = \Sigma$, yield affine invariant sampling schemes (see Appendix E for details).

Before proceeding to the important issue of selecting a practically useful choice of B, we observe the following important properties of our formulation: (i) the stochastic dynamical system (7) exactly preserves the target distribution and thus, if discretization error is well controlled, Metropolis correction is not likely to be needed, and (ii) the formulation, with appropriate choice of B, is affine invariant, even under discretization (see Appendix E), a property which ensures the stability of the method under change of coordinates. By contrast, we emphasize that schemes that modify S (instead of J) in (5) or that are based on a q-dependent normal distribution φ in (4) (e.g. within HMC as in [12]), cannot be made affine invariant in the same sense (Appendix E).

With the general stochastic quasi-Newton form in (7) as a template, one may consider many possible choices of B. Just as in optimization, in MCMC the question is not whether one should precondition, but rather how can one precondition in an affordable and effective way. Unfortunately, practical and effective quasi-Newton approaches for optimization do not have direct analogues in the sampling context leaving a substantial gap between un-preconditioned

methods and often impractical preconditioning approaches. In the next section we suggest an alternative strategy to fill this gap: using multiple copies of a simulation to incorporate local scaling information in the B matrix in (7).

3 Ensemble quasi-Newton (EQN) schemes

We next describe an efficient approach to the sampling problem in which information from an ensemble of walkers provides an estimate of the local covariance matrix. We consider a system of L walkers (independent copies evolving under the same dynamics) with state $x_i = (q_i, p_i)^{\mathrm{T}}$, where subscripts now indicate the walker index. Each walker has position q_i and momentum p_i for $i = 1, \cdots, L$, and with $Q = (q_1, q_2, \ldots, q_L)^{\mathrm{T}} \in \mathcal{D}^L$ and $P = (p_1, p_2, \ldots, p_L)^{\mathrm{T}} \in \mathbb{R}^{NL}$. We seek to sample the product measure $\bar{\pi}$ whose marginals give copies of the distribution of interest π:

$$\bar{\pi}(Q, P) = \prod_{i=1}^{L} \hat{\pi}(q_i, p_i), \qquad \int \bar{\pi}(Q, P) \mathrm{d}P = \prod_{i=1}^{L} \pi(q_i).$$

A simple sampling strategy is for each walker to sample $\bar{\pi}$ by evolving each x_i independently using an equation such as (2) or (5). Such a method scales perfectly in parallel when initial conditions are drawn from the target distribution, but no use is made of the local observed geometry or inter-walker information. Alternatively we may use the dynamics (6) to introduce walker information through the $B(q)$ matrix in order to precondition based on information from the other walkers. This preconditioning enters into the dynamics but not the distribution $\bar{\pi}$, ensuring it remains the product of disjoint distributions. Inserting the preconditioning matrix into the distribution in some form can have computational drawbacks associated to the communication of a large amount of information between steps.

Using L walkers, we have the global state $x = (Q, P)$ with $2NL$ total variables, with $B(Q)$ an $NL \times NL$ matrix. We will use $B(Q) = \mathrm{diag}(B_1(Q), B_2(Q), \ldots, B_L(Q))$ with each $B_i(Q) \in \mathbb{R}^{N \times N}$ so that the position and momentum (q_i, p_i) of walker i evolve according to (7) with $B(q)$ replaced by $B_i(Q)$. Note that the divergence and gradient terms in the equation for each walker are are taken with respect to the q_i variable.

With this quasi-Newton framework there are many potential choices for the B_i matrix, with $B_i = I_N$ reducing us to running L independent copies of underdamped Langevin dynamics. Before exploring the possibilities we remark that, in order to exploit parallelism, we will divide our L walkers into several groups of equal size. Walkers in the same group $g(i)$ as walker i will *not* appear in B_i so that the walkers in any single group can be advanced in parallel independently. We set $Q_{[i]} = \{q_j \,|\, g(j) \neq g(i)\}$ and let K be the common size of these sets. For example, if we have 16 cores available we may wish to use ten groups of 16 walkers (so $L = 160$ and $K = 144$). If walker j is designated as belonging to group 1, it evolves under the dynamics given in equation (7) but

the set $Q_{[j]}$ only includes walkers in groups $2, \cdots, 10$. We may then iterate over the groups of walkers sequentially, moving all the walkers in a particular group in parallel with the others. This compartmentalized parallelism is inspired by a similar approach in the emcee package [9] and provides high parallel efficiency.

One choice for the preconditioning matrix (not yet the one we employ) is to use the sample covariance of the ensemble

$$B_i(Q) = \sqrt{\operatorname{cov}(Q_{[i]})}. \tag{8}$$

Note that $\operatorname{div}(B_i(Q)^{\mathrm{T}}) \equiv 0$, simplifying the Metropolization of the scheme. In order for $B_i(Q)$ to be positive definite, we need at least N linearly independent walker positions, which at minimum requires that $L > N$.

With the choice of B_i in (8), the ensemble scheme applied to the density

$$\bar{\pi}_{A,v}(Q,P) = \prod \hat{\pi}_{A,v}(q_i,p_i) = \prod \hat{\pi}(Mq_i + v, p), \tag{9}$$

for some invertible matrix A and vector v, generates a sequence of vectors $(q_1^{(n)}, \ldots, q_L^{(n)}, p_1^{(n)}, \ldots, p_L^{(n)})$ with the property that the transformed sequence $(Aq_1^{(n)} + v, \ldots, Aq_L^{(n)} + v, p_1^{(n)}, \ldots, p_L^{(n)})$ has exactly the same distribution as the sequence generated by the ensemble scheme applied to $\bar{\pi}$ (see Appendix E). Just as choosing B as the square root of global covariance of π in (6) yields an affine invariant scheme, choosing the B_i as the square root of the ensemble covariance yields an affine invariant ensemble scheme. This affine invariance property suggests that ensemble schemes with B_i chosen as in (8) should perform well when the covariance of π has a large condition number. Note that a related choice in the context of an overdamped formulation and without the measure-correcting $\operatorname{div}(S)$ term (therefore reliant on an accept/reject step) appears in [16] and is shown to be affine invariant. An ensemble version of the HMC scheme using a mass matrix that seems to approximate the expectation of the Hessian matrix of $\log \pi$ appears in [28] but we have been unable to verify that the method as described preserves the target density (either exactly or up to a time-discretization error).

Using (8) in our ensemble schemes is problematic for several reasons. For high dimensional problems, the requirement that $L > N$ may render the memory demands of the methods prohibitive. This problem can be easily remedied by only approximating and rescaling in the space spanned by the eigenvectors corresponding to the largest eigenvalues of the ensemble covariance matrix. While such a scheme can be implemented in a reasonably efficient manner, we find that simply blending the sample covariance matrix with the identity via the choice

$$B_i(Q) = \sqrt{I_N + \mu \operatorname{cov}(Q_{[i]})}, \tag{10}$$

for some fixed parameter $\mu \geq 0$ is just as effective and much simpler. This choice allows $L \leq N$ but destroys affine invariance. On the other hand as demonstrated in Section (4), the method is still capable of dramatically alleviating scaling issues.

Having resolved the rank deficiency issue by moving to the choice of B_i in (10), one difficulty remains. As described in the previous section, for many problems we might expect that the global covariance of π is reasonably well scaled but that the sampling problem is still poorly scaled (the Hessian of $-\log\pi$ has large condition number). To address problems of this type, we define a localized covariance matrix that better approximates the Hessian at a point q_i while retaining full rank. We weight samples in the covariance matrix based on their distance (scaled by the global covariance) to a walker's current position, i.e. we use

$$B_i(Q) = \sqrt{I + \mu\,\mathrm{wcov}(Q_{[i]}, \omega_\lambda(Q_{[i]}, q_i))}, \tag{11}$$

for parameters μ, $\lambda > 0$, where now $\mathrm{wcov}(x, w)$ is a weighted covariance matrix of $K < L$ samples $q \in \mathbb{R}^{K\times N}$ with potentially unnormalized weights $w \in \mathbb{R}^K_+$

$$\mathrm{wcov}(q, w) = C \in \mathbb{R}^{N\times N}, \qquad C_{ij} = \sum_{k=1}^{K} \frac{w_k}{W}(q_{k,i}-\bar{q}_i)(q_{k,j}-\bar{q}_j), \qquad \bar{q}_i = \sum_{k=1}^{K} \frac{w_k}{W} q_{k,i}$$

with $W = \sum_k w_k$ and

$$\omega_\lambda(Q, q) = w \in \mathbb{R}^K_+, \qquad w_j = \exp\left(-\frac{\lambda}{2}\|Q_j - x\|^2_{\mathrm{cov}(Q)}\right),$$

where $\|y\|^2_X = y^{\mathrm{T}} X y$. Choosing $\lambda = 0$ reduces to (10), whereas a large value of λ gives more refined estimation of the local scaling properties of the system. The divergence term in (7) can be computed explicitly by computing partial derivatives of $B_i(q)$, making use of the formula for the derivative of the square root of a matrix: $\partial_i M(x) = M\Phi(M^{-1}(\partial_i(MM^T))M^{-T})$, where $\Phi(M) = \mathrm{lower}(M) + \mathrm{diag}(M)/2$. Note that the matrices $B_i B_i^{\mathrm{T}}$ for B_i in (11) are sums of the identity and L rank one matrices so that all manipulations involving B_i can be accomplished in linear cost in the dimension N. In Appendix B we detail a Metropolis–Hastings test that can be implemented (if needed) to correct any introduced bias. Because our ensemble scheme preserves π exactly when δt is small, one can also use the scheme absent of any Metropolis–Hastings test, improving the prospects for it to scale to very high dimension.

4 Numerical tests

We consider two numerical experiments to demonstrate the potential improvements that this method offers. A python package with example implementations of the code is available at [23].

4.1 Gaussian mixture model

We use the model presented in [4] focussing on fitting the distribution of a data vector y to univariate mixture model as the sum of three Gaussian distributions. The state vector is described by the means, precisions and weights of the three

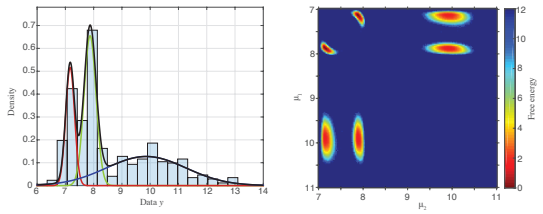

Figure 2: We plot a maximum likelihood state (left) with the three histograms colored as red, green and blue, with their sum in black, along with the original stamp data y as a histogram. The six modes due to label switching can be seen when looking at the free energy plot (right) in μ_1 and μ_2.

Gaussian distributions, denoted μ_i, λ_i and z_i respectively. Due to the sum of the weights equalling unity, this gives us $N = 8$ variables describing the mixture model. We also include a hyperparameter β that describes the rate parameter in the prior distribution on the precisions, giving a nine dimensional state. A full description of the problem is available in the Appendix C.

We consider the Hidalgo stamps benchmark dataset, studied in [20], as the data y with 485 datapoints. This example is well-suited to the local covariance approach we present above, due to the invariance of the likelihood under a permutation of components (the label-switching problem). Thus the system admits sets of $3! = 6$ equivalent modes, see Figure 2, each with a local scaling matrix that has the same eigenvalues with permuted eigenvectors.

As the walkers may initialize in the neighborhood of different local modes, using a 'global' preconditioning strategy would be sub-optimal as the best preconditioning matrix for the current position of a walker depends on which mode is closest to the walker. Instead, we use the covariance information from proximal walkers as in (11) to determine the optimal scaling.

We test the EQN scheme against the standard HMC scheme and a Metropolized version of Langevin dynamics. We used $L = 64$ walkers for the each scheme and compare the computed integrated autocorrelation times for an ensemble mean of quantities that vary slowly, shown in Table 1. The autocorrelation times are computed using the ACOR package [13].

Table 1: Computed autocorrelation times for slow variables

Scheme	$\min(z)$	$\max(\lambda)$	$\min(\mu)$	β
HMC	21495	42935	27452	7148
Langevin Dynamics	6825	13279	8384	4641
Ensemble Quasi-Newton	69	83	98	115

We consider all three methods as equivalent in cost, as they require the same number of evaluations of $\nabla \log(\pi)$ per step, and scale similarly with the size of

the data vector y. The EQN scheme is about 100 times more efficient compared to Langevin dynamics, and 350 times more efficient than HMC. We found that removing the divergence term in the EQN scheme had no significant impact on the results.

4.2 Log Gaussian Cox model

Figure 3: The synthetic observed intensity Y (left) and the true Gaussian field X (right).

To illustrate the method in a high-dimensional setting we compare results for inference in the Log Gaussian Cox point process as in [6]. We aim to infer the latent variable field from observation data. We make use of the RMHMC code template [11]. In the model, we discretize the unit square into a 32×32 grid, with the observed intensity in each cell denoted $Y_{i,j}$ and Gaussian field $X_{i,j}$. We use two hyperparameters σ^2 and β to govern the priors, making the dimensionality of the problem $N = 32^2 + 2 = 1026$ dimensions. Full details of the model are provided in Appendix D.

As the evaluation of the derivative of the likelihood is significantly cheaper with respect to the latent variables (tests showed computing the hyperparameter's derivatives to be about one hundred times slower), we employ a Gibbs scheme to first sample the latent variables using multiple steps, then perform one iteration for the hyperparameter distribution.

We generate synthetic test data Y, plotted in Figure 3, and compare the HMC and Langevin dynamics schemes to EQN (using 160 walkers) and the RMHMC scheme [12]. We additionally compare the results using the Langevin dynamics and EQN scheme without Metropolization, as the dynamics themselves sample π, and the Metropolis test only serves to remove discretization error (which is dominated by the sampling error in this example). RMHMC uses Hessian information to obtain scaling data for the distribution. This gives it a significant increase in cost, but improves the rate at which the sampler decorrelates. For the model, the RMHMC scheme requires about 2.2s per step whereas the other schemes require about 0.35s per step.

The EQN scheme significantly outperforms the other methods, with the slowest motion of the system (the β hyperparameter) decorrelating more rapidly than the HMC or Langevin schemes for approximately the same cost. The RMHMC scheme's requires significant extra computation, making it much less efficient than the standard HMC scheme in this example.

Table 2: Maximum autocorrelation times for each variable

Scheme	Latent variables	σ^2	β	Efficiency
HMC	800.7	1041.6	1318.7	1.0
RMHMC	2158.9	34.0	1502.0	0.15
Langevin Dynamics	405.1	140.6	435.3	3.5
Langevin Dynamics (non-Met)	81.6	20.5	136.5	11.2
Ensemble Quasi-Newton	71.9	49.2	239.5	5.4
Ensemble Quasi-Newton (non-Met)	64.4	8.8	47.8	26.8

5 Conclusion

We have presented a sampling algorithm that utilizes information from an ensemble of walkers to make more efficient moves through space, by discretizing a continuous ergodic quasi-Newton dynamics sampling the target distribution $\pi(x)$. The information from the other walkers can be introduced in several ways, and we give two examples using either local or global covariance information. The two forms of the B_i preconditioning matrix are then tested on benchmark test cases, where we see significant improvement compared to standard schemes.

The EQN scheme is cheap to implement, requiring no extra evaluations of $\nabla \log \pi(x)$ compared to schemes like MALA, and needing no higher derivative or memory terms. The scheme is also easily parallelizable, with communication between walkers being required infrequently. The dynamics (6) are novel in their approach to the introduction of the scaling information and we build on previous work using walkers running in parallel to provide a cheap alternative to Hessian data.

The full capabilities of the EQN method, in the context of complex data science challenges, remain to be explored. It is likely that more sophisticated choices of B_i are merited for particular types of applications. The propagation of an ensemble of walkers also suggests natural extensions of the method to sensitivity analysis and to estimation of the sampling error in the MCMC scheme. Also left to be explored is the estimation of the convergence rate as a function of the number of walkers, which may be possible for simplified model problems.

References

[1] C. H. Bennett. Mass tensor molecular dynamics. *Journal of Computational Physics*, 19(3):267–279, 1975.

[2] N. Bou-Rabee, A. Donev, and E. Vanden-Eijnden. Metropolis integration schemes for self-adjoint diffusions. *Multiscale Modeling & Simulation*, 12(2):781–831, 2014.

[3] C. J. T. Braak. A Markov chain Monte Carlo version of the genetic algorithm differential evolution: easy Bayesian computing for real parameter spaces. *Statistics and Computing*, 16(3):239–249, 2006.

[4] N. Chopin, T. Lelièvre, and G. Stoltz. Free energy methods for Bayesian inference: efficient exploration of univariate Gaussian mixture posteriors. *Statistics and Computing*, 22(4):897–916, 2012.

[5] J. A. Christen, C. Fox, et al. A general purpose sampling algorithm for continuous distributions (the t-walk). *Bayesian Analysis*, 5(2):263–281, 2010.

[6] O. F. Christensen, G. O. Roberts, and J. S. Rosenthal. Scaling limits for the transient phase of local Metropolis–Hastings algorithms. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):253–268, 2005.

[7] S. Duane, A. Kennedy, B. J. Pendleton, and D. Roweth. Hybrid Monte Carlo. *Physics Letters B*, 195(2):216 – 222, 1987.

[8] A. B. Duncan, T. Lelièvre, and G. A. Pavliotis. Variance reduction using non-reversible Langevin samplers. *Journal of Statistical Physics*, 163(3):457–491, 2016.

[9] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman. emcee: The MCMC hammer. *Publications of the Astronomical Society of the Pacific*, 125(925):306, 2013.

[10] W. Gilks, G. Roberts, and E. George. Adaptive direction sampling. *Journal of the Royal Statistical Society. Series D (The Statistician)*, 43(1):179–189, 1994.

[11] M. Girolami and B. Calderhead. Matlab code for the RMHMC scheme. `http://www.ucl.ac.uk/statistics/research/rmhmc`, 2011. [Online; accessed 01-Dec-2015].

[12] M. Girolami and B. Calderhead. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011.

[13] J. Goodman. ACOR package. `http://www.math.nyu.edu/faculty/goodman/software/`, 2009. [Online; accessed 01-Dec-2015].

[14] J. Goodman and A. Sokal. Multigrid Monte-Carlo method-conceptual foundations. *Phys Rev D*, 40(6):2035–2071, 1989.

[15] J. Goodman and J. Weare. Ensemble samplers with affine invariance. *Communications in applied mathematics and computational science*, 5(1):65–80, 2010.

[16] P. Greengard. An ensemblized Metropolized Langevin sampler. Master's thesis, NYU, 2015.

[17] H. Haario, E. Saksman, and J. Tamminen. An adaptive Metropolis algorithm. *Bernoulli*, 7(2):223–242, 2001.

[18] C.-R. Hwang, S.-Y. Hwang-Ma, and S.-J. Sheu. Accelerating gaussian diffusions. *The Annals of Applied Probability*, pages 897–913, 1993.

[19] C.-R. Hwang, S.-Y. Hwang-Ma, S.-J. Sheu, et al. Accelerating diffusions. *The Annals of Applied Probability*, 15(2):1433–1444, 2005.

[20] A. J. Izenman and C. J. Sommer. Philatelic mixtures and multimodal densities. *Journal of the American Statistical association*, 83(404):941–953, 1988.

[21] B. Leimkuhler, C. Matthews, and G. Stoltz. The computation of averages from equilibrium and nonequilibrium Langevin molecular dynamics. *IMA Journal of Numerical Analysis*, 2015.

[22] J. Martin, L. C. Wilcox, C. Burstedde, and O. Ghattas. A stochastic Newton MCMC method for large-scale statistical inverse problems with application to seismic inversion. *SIAM Journal on Scientific Computing*, 34(3):A1460–A1487, 2012.

[23] C. Matthews. Ensemble Quasi-Newton python package. `http://bitbucket.org/c_matthews/ensembleqn`, 2016. [Online; accessed 01-Jul-2016].

[24] G. Milstein and M. Tretyakov. *Stochastic Numerics for Mathematical Physics*. Springer, 2004.

[25] L. Rey-Bellet and K. Spiliopoulos. Irreversible langevin samplers and variance reduction: a large deviations approach. *Nonlinearity*, 28(7):2081, 2015.

[26] G. O. Roberts and J. S. Rosenthal. Coupling and ergodicity of adaptive Markov chain Monte Carlo algorithms. *Journal of applied probability*, pages 458–475, 2007.

[27] P. J. Rossky, J. D. Doll, and H. L. Friedman. Brownian dynamics as smart Monte Carlo simulation. *The Journal of Chemical Physics*, 69(10):4628–4633, 1978.

[28] Y. Zhang and C. A. Sutton. Quasi-newton methods for markov chain monte carlo. In *Advances in Neural Information Processing Systems*, pages 2393–2401, 2011.

Appendix

A Autocorrelation times for poorly conditioned problems

For comparisons of efficiency in MCMC methods, the integrated correlation time (IAT) τ is often used as a rough measure of the rate of convergence of the chain, giving the approximate time between independent samples. For a function of interest $f(x)$ with mean zero, its IAT τ_f is

$$\tau_f = 1 + \frac{2}{\mathrm{var}[f(x)]} \sum_{n=1}^{\infty} \mathrm{cov}[f(x^{(n)}), f(x^{(0)})],$$

where E_{x_0} denotes the expectation with respect to the initial conditions $x^{(0)}$. We shall consider sampling the distribution $\pi(x) \propto \exp(-x^T M^{-1} x/2)$ for some symmetric positive definite matrix M, in order to investigate the efficiency of the scheme (2)

$$x^{(n+1)} = x^{(n)} + \delta t \, \nabla \log(\pi(x^{(n)})) + \sqrt{2\delta t}\, \mathrm{R}^{(n)}.$$

As in Section VIII of [14] we will use the worst case IAT for functions of form $f(x) = v^T x$.

Suppose that $\lambda_{\max}$ is the largest eigenvalue of M^{-1} with corresponding eigenvector $v_{\max}$. Then if we choose $x^{(0)} = v_{\max}$

$$\mathrm{E}[x^{(n)}] = (1 - \delta t \lambda_{\max})^n v_{\max},$$

and hence we must choose δt such that $|1 - \delta t \lambda_{\max}| < 1$ to ensure convergence to the invariant mean vector zero. This gives the linear stability condition $\delta t < 2/\lambda_{\max}$.

We would expect that using a larger value of δt would lead to a more rapid decorrelation between samples. We can quantify this by computing the covariance between samples, where

$$\begin{aligned}\mathrm{cov}[x^{(n)}, x^{(0)}] &= (I - \delta t M^{-1})\mathrm{cov}[x^{(n-1)}, x^{(0)}] \\ &= (I - \delta t M^{-1})^n M\end{aligned}$$

from the definition of the update scheme. Similarly for a test function $f(x) = v^T x$ we have

$$\mathrm{cov}[f(x^{(n)}), f(x^{(0)})] = v^T (I - \delta t M^{-1})^n M v,$$

with integrated autocorrelation time

$$\tau_f = 1 + \frac{2}{v^T M v} v^T \left(\sum_{n=1}^{\infty} (I - \delta t M^{-1})^n \right) M v.$$

Assuming the linear stability condition is satisfied, we can rewrite this expression as

$$\tau_f = 1 + \frac{2}{\delta t \, v^T M v} v^T (M - \delta t I) M v = \frac{2 v^T M^2 v}{\delta t \, v^T M v} - 1. \tag{12}$$

Plugging in $v = v_{\min}$, the eigenvector corresponding to the minimum eigenvalue of M^{-1} (with eigenvalue $\lambda_{\min}$) we find that for the particular observable $f_*(x) = v_{\min}^T x$

$$\tau_{f_*} = \frac{2}{\delta t \, \lambda_{\min}} - 1,$$

and given the constraint $\delta t < 2/\lambda_{\max}$ this gives

$$\tau_{f_*} > \frac{\lambda_{\max}}{\lambda_{\min}} - 1,$$

so that even choosing the largest timestep possible, the rate of convergence will be slow whenever M has a wide range of eigenvalues.

B Metropolization of discretized scheme

In order to improve stability of the scheme, or to correct for numerical bias, we may seek to impose a Metropolis condition on the discretization of the dynamics (6). The discretization we use is given in (7), which we rewrite here for clarity:

$$p^{(n+1/4)} = p^{(n)} + \frac{\delta t}{2} F(q^{(n)}), \tag{13a}$$

$$q^{(n+1/2)} = q^{(n)} + \frac{\delta t}{2} B(q^{(n+1/2)}) p^{(n+1/4)} \tag{13b}$$

$$p^{(n+2/4)} = p^{(n+1/4)} + \frac{\delta t}{2} \mathrm{div}(B(q^{(n+1/2)})^{\mathrm{T}}) \tag{13c}$$

$$\hat{p}^{(n+2/4)} = \alpha p^{(n+2/4)} + \sqrt{1 - \alpha^2} \mathrm{R}^{(n)} \tag{13d}$$

$$p^{(n+3/4)} = \hat{p}^{(n+2/4)} + \frac{\delta t}{2} \mathrm{div}(B(q^{(n+1/2)})^{\mathrm{T}}) \tag{13e}$$

$$q^{(n+1)} = q^{(n+1/2)} + \frac{\delta t}{2} B(q^{(n+1/2)}) p^{(n+3/4)} \tag{13f}$$

$$p^{(n+1)} = p^{(n+3/4)} + \frac{\delta t}{2} F(q^{(n+1)}), \tag{13g}$$

with $\alpha = \exp(-\gamma\delta t)$, $\mathrm{R}^{(n)} \sim N(0, I)$ and $F(q) = B(q)^{\mathrm{T}} \nabla \log \pi(q)$. Note that the step in (13b) must be solved implicitly, likely requiring many evaluations of the matrix B. However, as this requires no communication between walkers and no evaluations of $\nabla \log \pi(q)$ we consider this a "cheap" operation.

The ratio of transition probabilities necessary for the acceptance rule is

$$\frac{T((q^{(n)}, p^{(n)}) \to (q^{(n+1)}, p^{(n+1)}))}{T((q^{(n+1)}, -p^{(n+1)}) \to (q^{(n)}, -p^{(n)}))} = \frac{f_\alpha(\mathrm{R}^{(n)})}{f_\alpha(\alpha\hat{p}^{(n+2/4)} - p^{(n+2/4)})} \frac{|I + \frac{1}{2}\delta t \nabla \otimes B(q^{(n+1/2)}) p^{(n+3/4)}|}{|I - \frac{1}{2}\delta t \nabla \otimes B(q^{(n+1/2)}) p^{(n+1/4)}|},$$

with derivatives taken with respect to q and where

$$f_\alpha(x) = \exp\left(-\frac{\|x\|^2}{2(1-\alpha^2)}\right).$$

In an efficient implementation, the calculation of the derivatives of $B(q)$ (needed for the transition probabilities and the divergence term) is only required once per step, between (13b) and (13c), while the evaluation of the $F(q)$ term is also once per step (after initialization) between lines (13f) and (13g). A single step can then be Metropolized using

- Set $(q^*, p^*) \leftarrow \mathrm{Upd}(q^{(n)}, p^{(n)})$
- Draw $u \sim U(0, 1)$
- Compute

$$U = \min\left(1, \frac{\hat{\pi}(q^*, p^*)}{\hat{\pi}(q^{(n)}, p^{(n)})} \frac{T((q^{(n)}, p^{(n)}) \to (q^*, p^*))}{T((q^*, -p^*) \to (q^{(n)}, -p^{(n)}))}\right)$$

- Then if $u < U$ we accept the move and set $(q^{(n+1)}, p^{(n+1)}) \leftarrow (q^*, p^*)$, otherwise we reject the move and flip the sign of the momentum, so set $(q^{(n+1)}, p^{(n+1)}) \leftarrow (q^{(n)}, -p^{(n)})$.

where the Upd function corresponds to a step of the discretization in (13). Similarly we can perform multiple steps and then accept/reject the trajectory by multiplying the associated transition probabilities. An example implementation in Python is available at [23]. For Metropolization of overdamped schemes such as (3), see [2].

C Details of the Gaussian mixture experiment

In this section we provide the details of the Gaussian mixture experiment for fitting the Hidalgo stamp data y to the density

$$\rho(x \,|\, \theta) = \sum_{k=1}^{3} z_k N(x \,|\, \mu_k,\, \lambda_k^{-1}),$$

where μ_k and λ_k are the center and precision of the component Gaussians, respectively, with weights $z_k > 0$ such that $\sum z_k = 1$. Let θ be the parameter/hyperparameter vector for this model. The target distribution for θ is $\pi(\theta) \propto p(\theta)\rho(y \,|\, \theta)$.

We use the prior distribution $p(\theta)$ such that for $k \in \{1, 2, 3\}$

$$\mu_k \sim N(m, \kappa^{-1}), \qquad \lambda_k \sim \mathrm{Gamma}(\alpha, \beta), \qquad (z_1, z_2, z_3) \sim \mathrm{Dirichlet}_3(1, 1, 1),$$

with hyperparameter $\beta \sim \text{Gamma}(g, h)$ and constants $m = \text{mean}(y)$, $r = \text{range}(y)$, $\kappa = 4/r^2$, $\alpha = 2$, $g = 0.2$, $h = 100g/(\alpha r^2)$.

We compare the standard HMC scheme, with a Metropolized Langevin dynamics scheme and the EQN scheme presented in the paper. For each of the schemes, we tweak the stepsize until the acceptance is on average about $75 - 80\%$. The HMC and Langevin schemes are run by taking 50 steps per single iteration, and using a Metropolis test on the obtained trajectory, while the EQN scheme takes 5 steps per iteration. The Langevin and EQN scheme used a friction of $\gamma = 0.01$.

All schemes used 64 walkers, which amounts to 64 independent runs for the HMC and Langevin schemes. The EQN run used the localized form of the covariance matrix, with $\mu = 100$ and $\lambda = 12$ in (11), however with the weighting kernel only using the Euclidean distance in the μ_i space, rather than the entirety of θ. We observed that increasing μ further reduced the acceptance probability significantly.

D Details of the log Gaussian Cox experiment

We run a larger experiment with 1024+2 total parameters to estimate. In the model, we discretize a unit square into a 32×32 grid, with the observed intensity in each cell denoted $Y_{i,j}$ and Gaussian field $X_{i,j}$.

The intensities are assumed to be conditionally independent and Poisson distributed with means $m\Lambda(i,j) = m\exp(X_{i,j})$ for latent intensity process $\Lambda(i,j)$ and $m = 1/32^2$. $X = \{X_{i,j}\}$ is a Gaussian process, where $x = \text{vec}(X)$ we have its mean $E(x) = \mu$ and covariance matrix

$$\Sigma_{(i,j),(i',j')} = \sigma^2 \exp(\delta(i,i',j,j')/32\beta), \quad \delta(i,i',j,j') = \sqrt{(i-i')^2 + (j-j')^2},$$

with parameters μ, σ^2 and β.

Our goal is to sample likelihoods for the latent variables X but also sample values for the hyperparameters β and σ^2, whose priors we assume are exponentially distributed.

We generate synthetic data Y from a field X, created using $\beta = 1/33$, $\sigma^2 = 1.91$ and $\mu = \log(126) - \sigma^2/2$. We aim to infer X from our synthetic Y, along with the hyperparameters used for the model, using the HMC, RMHMC, Langevin dynamics and ensemble quasi-Newton methods. We make use of the RMHMC template code for the problem, available at `http://www.ucl.ac.uk/statistics/research/rmhmc`.

In order to run multiple highly-resolved simulations, we use a 32×32 grid rather than the 64×64 grid used in the original version of the problem. This reduces the dimension of the latent variables from 4096 to 1024, which we still consider large enough for a significant test of the samplers' abilities, but this reduction allows us to run for longer and recover more accurate statistical information. However, the change in the model requires us to alter some parameters used in the RMHMC method, for example the timestep, in order to recover good efficiency. The timestep is increased until we reach 75% acceptance (though we do not claim our choice is optimal).

We implement all of the schemes in MATLAB, with each walker running on a single core of identical architecture. For all methods, one iteration uses 50 latent variable steps for each one hyperparameter step. The ergodic property of the Langevin-dynamics based schemes allows us to run without Metropolization, if we are willing to endure some discretization bias that is not removed with further sampling. We argue that in most practical cases, when using a sensible discretization parameter sampling error should always dominate the discretization bias.

For the ensemble quasi-Newton sampler, we run using 160 walkers using the global covariance formulation for B_i (effectively $\lambda = 0$). We use a Gibbs sampling approach to sample the latent variables and hyperparameters, with the schemes using a B_i for each partition of variables. For the hyperparameters we used $\mu = 1$, and for the latent variables we used $\mu = 7$. We use five groups of 32 walkers with the walkers within each group running in parallel for ten thousand steps before switching to the next group. The cost of this communication is negligible as it is done so infrequently.

E Affine invariance

We can extend our notion of affine invariance to the underdamped case, where the system state is $x = (q, p)$, with target distribution $\hat{\pi}(x) \propto \pi(q)\exp(-\|p\|^2/2)$. We consider affine transformations exclusively of the form $\hat{\psi}(x) = (\psi(q), p) = (Ax + v, p)$, defining the transformed distribution

$$\hat{\pi}_{\hat{\psi}} \propto \pi_\psi(q)\exp(-\|p\|^2/2), \qquad \pi_\psi(q) \propto \pi(\psi(q)),$$

for any invertible matrix A and fixed vector v. We may apply the discretization (13) to the density $\hat{\pi}_\psi$, using some scaling matrix $B_{\pi_\psi}(q)$ that we shall choose later, with $B_{\pi_\psi}(q) = B_\pi(\psi(q))$. The first computation (13a) for this density is

$$p^{(n+1/4)} = p^{(n)} + \frac{\delta t}{2} B_{\pi_\psi}(q)^{\mathrm{T}} \nabla \log \pi_\psi(q^{(n)}),$$

which is equivalent to, when writing $y = \psi(q)$,

$$p^{(n+1/4)} = p^{(n)} + \frac{\delta t}{2} B_\pi(y^{(n)})^{\mathrm{T}} A^{\mathrm{T}} \nabla \log \pi(y^{(n)}). \tag{14}$$

Whereas writing $q = A^{-1}(y - v)$ in line (13b) and multiplying by A we have

$$y^{(n+1/2)} = y^{(n)} + \frac{\delta t}{2} A B_\pi(y^{(n+1/2)}) p^{(n+1/4)}. \tag{15}$$

Finally, line (13c) becomes

$$p^{(n+2/4)} = p^{(n+1/4)} + \frac{\delta t}{2} g(y^{(n+1/2)}) \tag{16}$$

where $g(x) = \mathrm{div}(B_\pi(x)^{\mathrm{T}} A)$.

Line (13d) remains unchanged, given our choice of affine invariant map $\hat{\psi}$ does not affect the momentum p. Suppose now that we choose

$$B_{\pi_\psi}(q) B_{\pi_\psi}(q)^T = A^{-1} C(q) A^{-T},$$

for some symmetric positive definite matrix C, and thus $B_{\pi_\psi}(q) = A^{-1}\sqrt{C(q)}$. For this choice, the discretization steps in (14)-(16) become independent of the scale factor A, and hence and we obtain a scale-independent sampler when we use (13). One such choice is to use the square root of the (constant) covariance matrix $\mathrm{cov}_{\pi_\psi}(q)$ as the $B(q)$ matrix, as

$$B_{\pi_\psi} B_{\pi_\psi}^T = \mathrm{cov}_{\pi_\psi}(q) = \mathrm{cov}_\pi(\psi^{-1}(x)) = A^{-1}\mathrm{cov}_\pi(x) A^{-T}, \tag{17}$$

satisfying the invariance property as the covariance is automatically symmetric positive definite. Alternatively, one may choose B to be the inverse square root of the Hessian matrix of $\log \pi_\psi$, which shares this property (subject to some conditions on the Hessian).

We have shown that the discretization using this choice of B is affine invariant only up to affine transformations in the q variables. However, we may consider dynamics that are invariant under a transformation

$$\tilde{\psi}(q,p) = (Aq + v, A^{-T}p),$$

which acts in both components, where as before A_i is any invertible matrix and v_i is a vector. The dynamics we consider sample the distribution $\tilde{\pi}(q,p)$ ergodically, where

$$\tilde{\pi}(q,p) \propto \pi(q)\tilde{\varphi}(q,p), \qquad \tilde{\varphi}(q,p) = \exp(-p^T M_\pi^{-1}(q)p/2 - \log|M_\pi(q)|/2),$$

so that the marginal distribution of $\tilde{\pi}(q,p)$ in the momenta is $\pi(q)$. The mass matrix $M_\pi^{-1}(q)$ is a symmetric positive definite matrix that may be dependent on position q. For the choice of

$$J = \begin{bmatrix} 0 & -I \\ I & 0 \end{bmatrix} \qquad S = \begin{bmatrix} 0 & 0 \\ 0 & \gamma I \end{bmatrix}$$

in (1) we obtain the standard underdamped Langevin dynamics, and applying it to $\tilde{\pi}(q,p)$ with $\gamma = 1$ gives

$$\mathrm{d}q = M_\pi^{-1}(q)p\mathrm{d}t, \qquad \mathrm{d}p = \nabla \log \pi(q)\mathrm{d}t + \nabla \log \tilde{\varphi}(q,p)\mathrm{d}t - p\mathrm{d}t + \sqrt{2M_\pi(q)}\mathrm{d}W_t,$$

where the gradient ∇ is taken with respect to the position q. Now applying this dynamics to the transformed distribution instead

$$\tilde{\pi}_{\tilde{\psi}}(q,p) \propto \pi_\psi(q)\tilde{\varphi}_\psi(q,p) = \pi(Aq+v)\tilde{\varphi}(Aq+v, A^{-T}p)$$

we obtain

$$\mathrm{d}q = M_{\pi_\psi}^{-1}(q)p\,\mathrm{d}t, \tag{18}$$

$$\mathrm{d}p = \nabla \log \pi_\psi(q)\mathrm{d}t - \tfrac{1}{2}\nabla \log|M_{\pi_\psi}(q)|\mathrm{d}t - \tfrac{1}{2}\nabla(p^T M_{\pi_\psi}^{-1}(q)p)\mathrm{d}t - p\mathrm{d}t + \sqrt{2M_{\pi_\psi}(q)}\mathrm{d}W_t, \tag{19}$$

where ∇ denotes gradient with respect to the position coordinates. Changing variables $r = A^{-T}p$ gives

$$\mathrm{d}q = M_{\pi_\psi}^{-1}(q)A^T r\,\mathrm{d}t,$$

$$\mathrm{d}r = A^{-T}\nabla \log \pi_\psi(q)\mathrm{d}t - \tfrac{1}{2}A^{-T}\nabla \log|M_{\pi_\psi}(q)|\mathrm{d}t - \tfrac{1}{2}A^{-T}\nabla(r^T A M_{\pi_\psi}^{-1}(q)A^T r)\mathrm{d}t - r\mathrm{d}t + A^{-T}\sqrt{2M_{\pi_\psi}(q)}\mathrm{d}W_t.$$

Then writing $y = Aq + v$ we have $\pi_\psi(q) \propto \pi(y)$, and if we assume M is chosen such that

$$M_{\pi_\psi}^{-1}(q) = A^{-1}C(y)A^{-T} \tag{20}$$

for some symmetric positive definite matrix $C(q)$, then the dynamics become

$$\mathrm{d}y = C(y)r\,\mathrm{d}t,$$

$$\mathrm{d}r = \nabla \log \pi(y)\mathrm{d}t + \tfrac{1}{2}\nabla \log|C(y)|\mathrm{d}t - \tfrac{1}{2}\nabla(r^T C(y)r)\mathrm{d}t - r\mathrm{d}t + \sqrt{2C(y)}\mathrm{d}W_t,$$

eliminating the A scaling term in the dynamics and yielding affine invariance with respect to the transformations $\tilde{\psi}$.

Using the inverse covariance matrix as the mass is one such choice fo the M matrix, as similar to (17) we have

$$M_{\pi_\psi} = \mathrm{cov}_{\pi_\psi}(q) = \mathrm{cov}_\pi(\psi^{-1}(x)) = A^{-1}\mathrm{cov}_\pi(x)A^{-T}$$

satisfying (20). The inverse Hessian is another such matrix with this property.

This result applies directly to the RMHMC scheme [12] which uses a position dependent mass matrix, however it periodically redraws the momentum and sets $\gamma = 0$. If the inverse Hessian is used (assuming it remains symmetric positive definite), or another matrix such that (20) holds, then the resulting dynamics will be affine invariant under transformations $\tilde{\psi}(q,p) = (Aq + v, A^{-T}p)$.

Computationally there is no obvious benefit to including a p-scaling term in the affine transformation, as sampling $\pi(q)$ is the ultimate goal of our sampling, and the inclusion of momentum is to increase efficiency. However, the usage of (18) may become practically inefficient in the case of ensemble sampling, as the joint distribution will be

$$\bar{\pi}(Q,P) = \prod_{i=1}^{L} \pi(q_i)\exp(-p_i^T M_i^{-1}(Q)p/2 - \log|M_i(Q)|/2).$$

In the target joint distribution the walker positions are no longer independent. This complicates Metropolization of the scheme which will now require calculations involving all walkers when any one walker's position is changed. Additionally in each walker's dynamics (18), evaluating the $\nabla_{q_i}\bar{\pi}(Q,P)$ term requires computing derivatives of each of the L-many M_i matrices, causing a significant amount of additional computational overhead per step.

F Noisy estimation of the divergence term

The divergence of a positive definite matrix appears in many of the schemes we consider above, however the derivative is sometimes prohibitively computationally expensive to obtain, or infeasible to compute analytically. A Metropolization condition can be enforced to recover the correct sampling without this term, but we may instead approximate the divergence of a matrix $M(x)$ using a random update. Formally, for a small constant $\epsilon > 0$ and vector $R \in \mathbb{R}^N$ we have

$$Z = [M(x+\epsilon R) - M(x)]R = \epsilon\sum_{i=1}^{N} R_i\partial_i M(x)R + \epsilon^2\sum_{i=1}^{N}\sum_{j=1}^{N} R_iR_j\partial_i\partial_j M(x)R + O(\epsilon^3).$$

If $R \sim N(0,I)$ then taking the expectation of Z yields

$$(\mathrm{E}[Z])_i = \epsilon\sum_{j=1}^{N}\partial_i M_{i,j}(x) + O(\epsilon^3),$$

and so $\mathrm{E}[Z] = \epsilon\,\mathrm{div}(M(x)) + O(\epsilon^3)$. This gives a cheap "noisy" approximation of the divergence term that will introduce some small bias into the system (depending on the spectrum of $M(x)$).